

(19)



Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 0 899 660 A1

(12)

## EUROPEAN PATENT APPLICATION

(43) Date of publication:  
03.03.1999 Bulletin 1999/09

(51) Int Cl.<sup>6</sup>: G06F 9/46, G06F 13/10

(21) Application number: 98306916.2

(22) Date of filing: 28.08.1998

(84) Designated Contracting States:  
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU  
MC NL PT SE  
Designated Extension States:  
AL LT LV MK RO SI

(72) Inventor: Murata, Seiji, c/o Sony Corporation  
Shinagawa-ku, Tokyo 141 (JP)

(74) Representative: Turner, James Arthur et al  
D. Young & Co.,  
21 New Fetter Lane  
London EC4A 1DA (GB)

(30) Priority: 29.08.1997 JP 234977/97

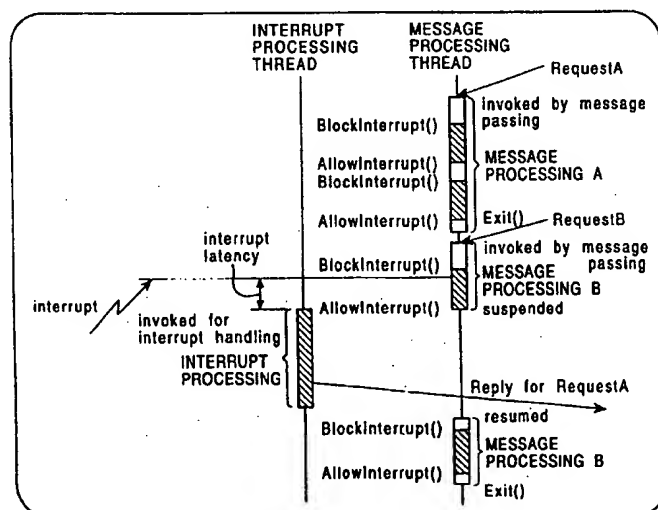
(71) Applicant: SONY CORPORATION  
Tokyo 141 (JP)

## (54) Data processing method and apparatus

(57) Each of object-oriented device drivers is described in terms of a multi-thread object which can allocate a message processing thread and an interrupt processing thread to be exclusively used for each interrupt. The device driver, when executing a processing based on a message received from another object, assigns the processing to the message processing thread. When an event has occurred requesting an interrupt to the device driver, the device driver executes the interrupt processing corresponding to the event in the interrupt processing thread corresponding to the event. In

case that the corresponding processing thread is busy due to execution of another interrupt processing in response to an earlier interrupt when the event has occurred, the interrupt processing corresponding to the event is held without being executed regardless of the state of the message processing thread, and is executed only after completion of the interrupt processing for the earlier interrupt that is under execution in the interrupt processing thread. Interrupt latency is shortened without affecting advantages brought about by the use of the object-oriented device drivers.

FIG. 4



EP 0 899 660 A1

## Description

[0001] This invention relates to data processing methods, recording media and data processing apparatus.

[0002] In general, device drivers that are programs for controlling various hardware devices are implemented as part of an operating system, as illustrated in Fig. 8. More specifically, the operating system, in response to a request for using a certain device driver given by a client or to a hardware interrupt, retrieves a device driver table to find the target device driver complying with the request or the interrupt from among device drivers registered in the device driver table. The operating system then calls the found device driver by means of a function call and executes the called device driver.

[0003] The client's request is given by means of a system call, while the hardware interrupt is made by invoking a kernel thread of the operating system.

[0004] This known system has the following problems.

(1) Adaptability of the operating system is limited and replacement of a device driver is not easy because the device driver is implemented as part of a kernel of the operating system.

(2) Any error in a device driver affects the whole operating system because the device driver is executed with the same authority as the kernel of the operating system.

(3) Interrupt and other controls are performed on independent device drivers. The programmers who form the device drivers, therefore, are required to have knowledge of the whole system and to consider the influence of each device driver on other parts of the system.

(4) The style of programming of the device driver is significantly different from that for application programs because the device driver is implemented as part of the kernel of the operating system. Therefore, it is not easy for a programmer to form a device driver unless he is familiar with and well trained in the programming style which is used in describing device drivers.

[0005] In order to overcome these problems, a method has been known in which device drivers are described by using an object-oriented technique. When the object-oriented technique is used, the device drivers are formed as parallel objects outside the operating system, in the same way as that for application programs, as shown in Fig. 9 by way of example.

[0006] The term "parallel objects" is used in this specification to mean program functions in the form of object modules which are objective entities to be handled by the operating system. Thus, the parallel objects are the units to be handled or processed. For instance, scheduling by the operating system is performed on a parallel-object basis. Communication also is executed on the

units of parallel objects. In other words, the objects constitute units for operations such as resource management and exclusive execution. The use of such objects enables a programmer to form an object without requiring knowledge of contents of other objects. Thus, the contents of each object can be determined without taking into account factors such as exclusive control.

[0007] The described method relying upon the object-oriented technique has the following two major features.

(1) A request from a client to a device driver is sent in the form of a message from the client to the device driver described as a parallel object, as indicated by "request" in Fig. 9. Similarly, an interrupt request is sent, as indicated by "interrupt" in the Fig. 9, in the form of a message from the operating system to the device driver described as a parallel object.

(2) Each device driver runs with a single thread, and the control of the interrupt mask is performed on independent device drivers. Thus, when a device driver is operating, interrupts to be handled by this device driver are not accepted. This relieves the programmers from the burden of considering exclusive control and other controls to cope with an interrupt when they describe each device driver.

[0008] By virtue of the two major features set forth above, device drivers described as parallel objects can be formed in the same way as that for application programs. In addition, replacement of device drivers can easily be carried out because the device drivers are implemented outside the operating system. It is also to be appreciated that the stability of operation of the operating system is greatly improved because an error occurring in one device driver does not affect the entire operating system.

[0009] Thus, the described problems encountered with the device drivers implemented as part of an operating system in the manner shown in Fig. 8 can be overcome by the technique shown in Fig. 9 in which the device drivers are described as parallel objects similar to application programs.

[0010] The technique shown in Fig. 9, however, is still unsatisfactory in that a running device driver does not accept any interrupt. A concept termed as "interrupt mask" is used to express the state of the device driver as to whether or not the device driver accepts an interrupt. For instance, the state of the device driver which rejects an interrupt is expressed as "the interrupt mask is closed". Thus, an expression "interrupt mask is open" means that the device driver is in a state ready to accept the interrupt.

[0011] In the technique shown in Fig. 9, the interrupt mask is closed when the device driver is running. Consequently, interrupt latency is prolonged when the processing which is being executed by the device driver includes time-consuming work such as copying of data.

[0012] Thus, interrupts cannot be accepted when the device driver is running because each device driver runs with a single thread. Hitherto, the only solution to this problem was to shorten the method of each device driver. It has been therefore difficult to describe a device driver which poses a severe restriction of interrupt latency when describing the device drivers in the form of parallel objects.

[0013] Various respective aspects of the invention are defined in the appended claims.

[0014] Embodiments of the present invention can provide a data processing method based on object-oriented device drivers and capable of shortening interrupt latency without impairing the advantage offered by the use of the device drivers.

[0015] Embodiments of the invention can provide a recording medium storing a program which implements the data processing method.

[0016] Embodiments of the invention can provide a data processing apparatus incorporating such a recording medium.

[0017] According to embodiments of the present invention, there is provided a data processing method in which a device driver for driving a hardware device is described in terms of a multi-thread object capable of allocating a message processing thread and interrupt processing threads to be exclusively used for respective interrupts. The device driver assigns to the message processing thread a processing based on a message received by the device driver from another object thereby executing the processing in the message processing thread. When an event has occurred requesting an interrupt to the device driver, the device driver executes an interrupt processing corresponding to the event in a processing thread corresponding to the event, wherein, in case that the corresponding interrupt processing thread is busy due to execution of another interrupt processing in response to another interrupt when the event has occurred, the interrupt processing corresponding to the event is held without being executed regardless of the state of the message processing thread and is executed only after completion of the interrupt processing which is under execution in the corresponding interrupt processing thread.

[0018] For the purpose of achieving synchronization to enable an exclusive control as required in, for example, making access to common data, it is preferred that an order of priorities is set between processing in the message processing thread and processing in the interrupt processing thread. In case that a processing given a higher priority than the interrupt processing is being executed in the message processing thread when an interrupt is received by the device driver, execution of the interrupt processing corresponding to the received interrupt is delayed until the processing in the message processing thread is finished or until the setting of the priority order is changed so as to give higher priority to the interrupt processing than to the processing in the

message processing thread. However, in case that the processing given the higher priority than the interrupt processing is not being executed in the message processing thread when the interrupt is received by the device driver, the interrupt processing is executed in the interrupt processing thread without delay after the receipt of the interrupt.

[0019] Preferably, an interrupt processing is divided into a plurality of portions having different levels of priority, and a portion of the interrupt processing of a higher priority is executed by being assigned to the interrupt processing thread while a portion of a lower priority is executed by being assigned to the message processing thread. This serves to shorten the time over which the interrupt processing is occupied, thus reducing interrupt latency.

[0020] In the data processing method of embodiments of the invention as set forth above, each device driver is described as a multi-thread object. Therefore, any interrupt can be accepted by the interrupt processing thread even when a processing is under execution by the device driver, provided that such a processing has been assigned to and being executed in the message thread.

[0021] In accordance with another embodiment of the present invention, there is provided a computer-readable recording medium storing a device driver which is a program for driving a hardware device. The device driver is described in terms of a multi-thread object capable of allocating a message processing thread and interrupt processing threads to be exclusively used for respective interrupt processings. The device driver assigns to the message processing thread a processing based on a message received by the device driver from another object thereby executing the processing in the message processing thread. When an event has occurred requesting an interrupt to the device driver, the device driver executes the interrupt processing corresponding to the event in a processing thread corresponding to the event, wherein, in case that the corresponding interrupt processing thread is busy due to execution of another interrupt processing in response to another interrupt when the event has occurred, the interrupt processing corresponding to the event is held without being executed regardless of the state of the message processing thread and is executed only after completion of the another interrupt processing in the corresponding interrupt processing thread.

[0022] Preferably, the device driver is arranged to allow setting of an order of priorities between processing in the message processing thread and processing in the interrupt processing thread. In case that a processing given a higher priority than the interrupt processing is being executed in the message processing thread when an interrupt is received by the device driver, execution of the interrupt processing corresponding to the received interrupt is delayed until the processing in the message processing thread is finished or until the set-

ting of the priority order is changed so as to give higher priority to the interrupt processing than to the processing in the message processing thread, whereas, in case that the processing given the higher priority than the interrupt processing is not being executed in the message processing thread when the interrupt is received by the device driver, the interrupt processing is executed in the interrupt processing thread without delay after the receipt of the interrupt.

[0023] The device driver may divide an interrupt processing into a plurality of portions having different levels of priority. In such a case, a portion of the interrupt processing of a higher priority is executed by being assigned to the interrupt processing thread while a portion of a lower priority is executed by being assigned to the message processing thread.

[0024] In the recording medium of embodiments of the invention as set forth above, each device driver is described as a multi-thread object. Therefore, any interrupt can be accepted by the interrupt processing thread even when a processing is under execution by the device driver, provided that such a processing has been assigned to and being executed in the message thread.

[0025] In accordance with still another embodiment of the present invention, there is provided a data processing apparatus, comprising a recording medium storing a device driver which is a program for driving a hardware device and which is readable by the apparatus from the recording medium. The device driver is described in terms of a multi-thread object capable of allocating a message processing thread and at least one interrupt processing thread. The device driver assigns to the message processing thread a processing based on a message received by the device driver from another object thereby executing the processing in the message processing thread. When an event has occurred requesting an interrupt to the device driver, the device driver executes the interrupt processing corresponding to the event in a processing thread corresponding to the event, wherein, in case that the corresponding processing thread is busy due to execution of another interrupt processing in response to another interrupt when the event has occurred, the interrupt processing corresponding to the event is held without being executed regardless of the state of the message processing thread, and is executed only after completion of the interrupt processing in the corresponding interrupt processing thread.

[0026] Preferably, the device driver is arranged to allow setting of an order of priorities between processing in the message processing thread and processing in the interrupt processing thread. In case that a processing given a higher priority than the interrupt processing is being executed in the message processing thread when an interrupt is received by the device driver, execution of the interrupt processing corresponding to the received interrupt is delayed until the processing in the message processing thread is finished or until the set-

ting of the priority order is changed so as to give higher priority to the interrupt processing than to the processing in the message processing thread, whereas, in case that the processing given the higher priority than the interrupt processing is not being executed in the message processing thread when the interrupt is received by the device driver, the interrupt processing is executed in the interrupt processing thread without delay after the receipt of the interrupt.

[0027] The device driver may divide an interrupt processing into a plurality of portions having different levels of priority. In such a case, a portion of the interrupt processing of a higher priority is executed by being assigned to the interrupt processing thread while a portion of a lower priority is executed by being assigned to the message processing thread.

[0028] In the data processing method of embodiments of the invention as set forth above, each device driver is described as a multi-thread object. Therefore, any interrupt can be accepted by the interrupt processing thread even when a processing is under execution by the device driver, provided that such a processing has been assigned to and being executed in the message thread.

[0029] The invention will now be described by way of example with reference to the accompanying drawings, throughout which like parts are referred to by like references, and in which:

Fig. 1 schematically shows the construction of a TV system incorporating an embodiment of the present invention;

Fig. 2 shows a device driver execution model;

Fig. 3 shows, by way of example, an operation for controlling an interrupt mask of a device driver in accordance with a known method under a type 1 programming style;

Fig. 4 shows, by way of example, an operation for controlling an interrupt mask of a device driver in accordance with the present embodiment under a type 2 programming style;

Fig. 5 shows, by way of example, an operation for controlling an interrupt mask of a device driver in accordance with the present embodiment under the type 2 programming style;

Fig. 6 shows, by way of example, an operation for controlling an interrupt mask of a device driver in accordance with the present embodiment under a type 3 programming style;

Fig. 7 shows, by way of example, an operation for controlling an interrupt mask of a device driver in accordance with the present embodiment under the type 3 of programming style;

Fig. 8 is an illustration of a conventional device driver programming in which device drivers are implemented as part of an operating system; and

Fig. 9 is an illustration of an object-oriented device driver programming in which device drivers are im-

plemented as parallel objects.

[0030] Fig. 1 shows, by way of example, a hardware architecture incorporating embodiments of the present invention. Although a TV system is specifically mentioned in the following description, it is to be understood that the present techniques can also be applied to a variety of types of data processing apparatuses of the kind in which various hardware components are controlled by device drivers. For instance, the present invention can effectively be employed in audiovisual systems, i. e., so-called AV systems, office machines, computers, and so forth, as well as in the television system which will now be described.

[0031] Referring to Fig. 1, the TV system as an embodiment of the data processing apparatus receives signals from a station via air (antenna) or through a cable, and displays an image based on the received signal on a CRT or a liquid crystal display device while outputting acoustic signals from a speaker.

[0032] This TV system has, besides ordinary TV functions, a function for receiving programs and data which are given externally. More specifically, the TV system as shown in Fig. 1 has a TV function unit 3 connected to a BUS 2 through a BUS/IO bridge 1, a processor 5 connected to the BUS 2 through a BUS/MEMORY bridge 4, a ROM (Read Only Memory) 6 and a RAM (Random Access Memory) 7 which also are connected to the processor 5 through the BUS/MEMORY bridge 4, an operation panel 8 connected to the BUS 2, and an external storage device 9 and a communication device 10 which are also connected to the BUS 2.

[0033] The TV function unit 3 has a function for generating images and voices based on signals received via air or a cable, and is connected to the BUS 2 through the BUS/IO bridge 1 for communication with other devices of the system.

[0034] The processor 5 controls operations of various devices constituting the TV system. To this end, the processor 5 is connected to the BUS 2 via the BUS/MEMORY bridge 4. To the processor 5 are also connected the ROM 6 and the RAM 7 via the BUS/MEMORY bridge 4.

[0035] The ROM 6 stores an operating system, device drivers and application programs which implement the controls to be performed by the processor 5. In this embodiment, the operating system is an object-oriented operating system. The device drivers and the application programs operate as parallel objects.

[0036] The RAM 7 is used as a work area for the processor 5. Thus, the processor 5 uses the RAM 7 as the work area when executing the operating system, device drivers and the application programs, thereby controlling the operations of various devices constituting the TV system.

[0037] The operation panel 8 serves as an input device which receives operational instructions input by the user, such as instructions for channel change-over or

volume control. Practically, the operation panel 8 has an input device such as a console having a plurality of buttons, a pointing device such as a mouse, and so on. Signals input through the operation panel 8 are delivered to the processor 5 via the BUS 2 and the BUS/MEMORY bridge 4. Based on the signals input through the operation panel 8, the processor 5 performs various computations to control the component devices.

[0038] The external storage device 9, which may be a hard disk device, stores image data, control data and application programs down-loaded from an external system via the communication device 10. Thus, the communication device 10 serves as an I/O unit which undertakes signal exchanges with the external system, and comprises, for example, a MODEM or a terminal adapter.

[0039] The TV system not only performs ordinary TV functions provided by the TV function unit 3 but also has various other functions. For instance, the TV system can receive programs and data received from external systems via the communication device 10. This function enables updating of versions of the operating system and the device drivers simply by receiving new software modules from an external network via the communication device 10.

[0040] This TV system executes, under the control of the processor 5, the operating system and the device drivers stored in the ROM 6, and allows application programs stored in the ROM 6 or the external storage device 9 to be executed on the operating system, thereby controlling the component devices. Thus, the ROM 6 is used as a computer-readable storage medium which stores the operating system and the device drivers. The operating system and the device drivers, however, may be stored in the RAM 7 or in the external storage device 9. Storage in the RAM 7 or the external storage device 9 is preferred when rewriting of the operating system and device drivers is necessary.

## 2. Software environment

[0041] A description will now be made as to the software environment of the TV system.

### 2-1 Operating system

[0042] The operating system used in this TV system is of the type known as an "object-oriented" operating system. Various application programs are executed on this operating system. These application programs include, for example, an application program for enabling the TV function unit 3 to display moving pictures, and an application program which implements a graphical user interface (GUI) for enabling the user to control the system through the operation panel 8.

[0043] This operating system can simultaneously provide execution environments for a plurality of programs. In the following description, each program execution en-

vironment will be referred to as a "meta-space". More specifically, the following three types of meta-space are provided by the operating system: a meta-space for executing the device driver (this meta-space will be referred to as "mDrive meta-space", hereinafter), a meta-space for enabling execution of a procedural application program such as an application program for enabling the TV function unit 3 to display moving pictures, and a meta-space for executing an object-oriented application program for implementing, for example, the GUI which enables the user to interact with the system.

[0044] In order to deal with interrupts, the mDrive meta-space provides a function for implementing various means such as means for registering an interrupt handler which is an interrupt handling object, means for invoking the handler in response to an interrupt, and means for controlling an interrupt mask. The mDrive meta-space also provides a basic message communication function and memory managing function which enable the device drivers to operate as parallel objects.

[0045] Each device driver operates in accordance with the function provided through an API (Application Program Interface) of the mDrive meta-space. Thus, when one of the device drivers is executed to implement a control in such a manner as to interfere with the operation of another device driver, e.g., control of an interrupt mask, the control is necessarily conducted via the API of the operating system, while direct mutual control between these device drivers is prohibited. This eliminates any risk of malfunction attributable to direct mutual control between the device drivers, thus ensuring stable operation of the device drivers.

[0046] The feature that the control of each device driver is conducted necessarily through the API of the operating system enables each device driver to operate as an independent module. It is therefore possible to dynamically alter each device driver without substantial difficulty. This permits an easy dynamic updating of each device driver, simply by down-loading a program module through a network. 2-2 Basic execution model of device driver

[0047] A description will now be given of a basic execution model of the device driver which is executable on the above-described operating system.

[0048] Fig. 2 shows, by way of example, a basic execution model of the device driver. The device driver is triggered by a message from a client, controls the hardware in accordance with the message, and terminates its operation upon completion of the control of the hardware. In this case, an application program that uses the device driver functions as the client. In the following description, the processing conducted by the device driver based on the message from the client will be referred to as "message processing".

[0049] In the event that an interrupt request is given to the device driver from the hardware, the device driver performs interrupt processing in accordance with the interrupt request.

[0050] More specifically, when an event has occurred to request an interrupt to the device driver, the device driver interrupts a thread of another application program under execution, if any. Then, the interrupt handler is triggered to execute the interrupt processing, while performing message exchange with the client as required. The thread of the application program that has been interrupted is then resumed or, depending on the kind of the interrupt, a client's method is triggered, when the interrupt processing performed by the interrupt handler is finished.

[0051] The interrupt handler is dynamically registered for each device driver when an initializing method is executed on the device driver. When an event has occurred that requests an interrupt to the device driver, a message requesting the interrupt is sent to the interrupt handler as a specific asynchronous message. Upon receipt of the request, the interrupt handler is triggered to execute the interrupt processing as described. The method of the interrupt handler is described as a method which is triggerable by message communication, as in the case of methods of other objects. When an interrupt handler is dynamically registered for each driver through the execution of the initializing method on each device driver, information concerning the message for each interrupt handler is registered in the mDrive meta-space.

### 2-3 Interrupt mask of device driver

[0052] The interrupt processing employs an interrupt mask of the device driver. The interrupt mask is controlled in the following manner. The description will proceed based on the following scenario. A message "RequestA" is input to a device driver from an object which functions as the client for the device driver. This object will be referred to as a "client object", hereinafter. Upon receipt of the message "RequestA", the device driver executes a predetermined message processing (referred to as "message processing A", hereinafter) for controlling the hardware, in accordance with the message "RequestA".

(2) When the processing conducted by the device driver is completed, the device driver releases the thread in which the message processing A has been executed, regardless of the progress of the processing performed by the hardware. In this state, the device driver delays replying to the message "RequestA" until an interrupt indicating completion of the processing performed by the hardware is received from the hardware.

(3) Then, a message "RequestB" is input to the device driver from another object (referred to as an "object B") which also functions as a client to the device driver. Upon receipt of the message "RequestB", the device driver executes a predetermined message processing (referred to as "message processing B", hereinafter) for controlling the

hardware, in accordance with the message "RequestB".

(4) It is assumed here that the processing that has been under execution by the hardware in accordance with the message "RequestA" is completed during the execution of the message processing B. This situation is an event that requests an interrupt to the device driver. Therefore, an "interruption message" is input from the hardware to the device driver via the operating system requesting execution of a predetermined interrupt processing which is bound to be executed in response to completion of the processing by the hardware.

(5) The device driver operates the interrupt handler to execute the predetermined interrupt processing in accordance with the received interrupt message and, upon completion of the interrupt processing, sends the pending reply to the message "RequestA" to the client object A.

[0053] The term "interrupt latency" is used in this specification to mean the length of time from the occurrence of an interrupt to the triggering of the interrupt handler responding to the interrupt. The present technique effectively shortens the interrupt latency. The interrupt processing cannot be executed when the device driver is in a condition which rejects any interrupt request, i.e., when the interrupt mask is closed. In general, therefore, the maximum interrupt latency coincides with the maximum masking period in which the mask is kept closed.

#### 2-3-1 Interrupt mask under known art

[0054] Before turning to the description of the preferred embodiment, a description will be given as to the manner in which an interrupt mask is controlled in a known device driver.

[0055] Fig. 3 illustrates, by way of example, an interrupt mask control executed in a known art in which the device driver is described as a parallel object. The device driver is described as a single-thread object. Control of the interrupt mask is performed by the operating system through switching of the thread, i.e., by context switching.

[0056] More specifically, the interrupt mask is "open" when the device driver has not been started, i.e., when the thread of the device driver is empty. Conversely, the interrupt mask is "closed" when the device driver has been started so that the thread of the device driver is active.

[0057] In other words, in the conventional device driver, the interrupt mask is in the closed state not only when an interrupt processing is being executed in response to an interrupt from the hardware but also when message processing in response to a message from another object, e.g., the message processing A or the message processing B in the above-described scenario.

[0058] Thus, as will be seen from Fig. 3, the interrupt

which was triggered during execution of the message processing B is held, without being executed, until the message processing B is completed. The interrupt handler is started to execute the interrupt processing corresponding to the interrupt message only after the completion of the message processing B.

[0059] In the conventional device driver, the interrupt mask is kept closed when a message processing is being executed in response to a message from another object. This naturally leads to prolongation of the interrupt latency. This problem is serious particularly when the message processing which is being executed takes a long time. In such a case, the interrupt latency is unduly extended so as to hamper the expected control of the hardware which poses a very severe restriction in regard to interrupt latency, resulting in an interrupt failure. Thus, the conventional device driver inherently has a risk that it may fail to correctly control the hardware due to excessive prolongation of the interrupt latency.

[0060] The only relief to this problem in the conventional device driver is to minimize the size of description of each method contained in the device driver or to separate as much as possible any portion that is not essential for the hardware control from the device driver itself. Consequently, there is a practical limit in the shortening of the interrupt latency in the conventional device driver.

[0061] In the following description, the programming style in which the device driver is described as a single-thread object, as in the above-described example, will be referred to as "type 1".

#### 2-3-2 Interrupt mask in accordance with the embodiments

##### (Version 1)

[0062] A description will now be given of an embodiment of the present invention. In this embodiment, a device driver is described as a multi-thread object. This programming style will be referred to as "type 2" in the following description.

[0063] Fig. 4 shows an execution model which follows the scenario described before, in particular an interrupt mask control performed in this embodiment. As will be seen from Fig. 4, the device driver is described as a multi-thread object. More specifically, a message processing thread and an interrupt processing thread are allocated to the device driver.

[0064] The interrupt processing thread is a thread which executes interrupt processing. When an interrupt handler is registered, the interrupt processing thread is allocated as a thread which is to be used exclusively for the newly registered interrupt handler. The other thread, i.e., the message processing thread, is intended to execute message processing. Only one message processing thread is allocated to each device driver. Thus, the illustrated embodiment employs different threads for the interrupt processing and the message

processing.

[0065] The interrupt mask is closed only when the interrupt handler is operating. It is assumed here that an event has occurred which requests an interrupt to the device driver. This request is held without being accepted if the interrupt processing thread of the device driver is busy due to execution of an interrupt processing responding to another interrupt request. Thus, the interrupt is held without being executed until the interrupt processing thread becomes vacant. However, when the interrupt processing thread is not busy, i.e., when it is not executing any previous interrupt processing, the interrupt request is accepted by the device driver.

[0066] This permits acceptance of an interrupt even when message processing is being executed in the message processing thread, thus offering shorter interrupt latency. Thus, the illustrated embodiment effectively shortens the interrupt latency without impairing the advantage which is brought about when the device driver is described as a parallel object.

[0067] When an interrupt has occurred requesting control of the hardware while message processing is being executed, it is necessary to achieve synchronization between the message processing and the interrupt processing, since the interrupt handler is not allowed to operate when the message processing is being executed. In other words, it is necessary to execute an exclusive control in order that the interrupt processing by the interrupt handler is correctly performed even when the interrupt has occurred during the execution of message processing.

[0068] An exclusive control is therefore performed in such a manner that a message processing thread controls the state of the interrupt mask by means of a predetermined API (Application Program Interface) so as to suppress starting of the interrupt handler in a period covering the period necessitating such an exclusive control, thereby ensuring that the interrupt processing by the interrupt handler is executed without fail.

[0069] It is to be noted that the interrupt masks controlled through the API implementing such an exclusive control are used only for the interrupts which are to be handled by the independent device drivers. Therefore, control of the interrupt mask of a certain device driver does not disturb operations of other device drivers.

[0070] A detailed description will now be given of the execution model of the type 2 programming style, with specific reference to Fig. 4.

[0071] Referring to Fig. 4, a method "BlockInterrupt()" and a method "AllowInterrupt()" have been provided as predetermined APIs by the mDrive meta-spaces. The method "BlockInterrupt()" controls an interrupt mask such that, even when an interrupt has occurred during message processing, the message processing is executed preferentially while blocking the interrupt, whereas the method "AllowInterrupt()" sets up such a procedure that allows an interrupt to be executed preferentially while suspending the message processing when

the interrupt has occurred during the execution of message processing.

[0072] Referring further to Fig. 4, a message "RequestA" is input to the device driver so that a message processing A is started in response to the message "RequestA" (invoked by message passing). The message processing A is executed with the message processing thread. In this message processing A, each period, shown by hatching, between a moment at which the method "BlockInterrupt()" is called and a moment at which the method "AllowInterrupt()" is called, permits the message processing to be preferentially performed while the interrupt is blocked. Therefore, any interrupt received during this period is held without being executed.

[0073] The device driver, upon completion of the processing to be executed by the device driver, executes a method "Exit()" thereby opening the message processing thread, regardless of the state or progress of the processing performed by the hardware. The device driver delays replying to the message "RequestA" until an interrupt indicating completion of the processing on the hardware is received from the hardware.

[0074] Then, another message "RequestB" is input to the device driver, so that a message processing B is started in response to the message "RequestB" (invoked by message passing). The message processing B is executed with the message processing thread. In this message processing B too, each hatched period, between a moment at which the method "BlockInterrupt()" is called and a moment at which the method "AllowInterrupt()" is called, is the period that permits the message processing to be preferentially performed while the interrupt is blocked. Therefore, any interrupt received during this period is held without being executed.

[0075] It is assumed here that the processing which has been under execution by the hardware in response to the message "RequestA" is completed while the message processing B is being executed. This is an event that requests an interrupt to the device driver. Thus, an interrupt message "Interrupt" is input from the hardware to the device driver through the operating system requesting execution of a predetermined interrupt processing that follows the completion of the processing by the hardware.

[0076] Referring to Fig. 4, however, the moment of occurrence of this interrupt falls within the period between the moment of calling the method "BlockInterrupt()" in response to the message "RequestB" and the moment of calling the method "AllowInterrupt()". Therefore, the message "Interrupt" is held without being executed. Thus, the procedure which gives priority to the message processing performed by the message processing than to the interruption processing is valid when the interrupt is received by the device driver, so that the interrupt processing is held waiting for execution.

[0077] Subsequently, the method "AllowInterrupt()" is called, so that the message processing B is suspended.



Thus, a context switching is performed from the message processing thread to the interrupt processing thread, whereby the interrupt processing is started in response to the interrupt message "Interrupt" (Invoked for interrupt handling). The interrupt processing is executed with the interrupt processing thread. Thus, the interrupt latency in this embodiment corresponds to the period between the moment at which the interrupt has occurred and the moment at which the method "AllowInterrupt()" is called during the processing B so as to allow the interrupt. The interrupt is then executed after expiration of this period of interrupt latency.

[0078] Thus, the device driver enables the interrupt handler to execute a predetermined interrupt responding to the interrupt message "Interrupt", and sends a reply (Reply for RequestA) to the message "RequestA" indicating the completion of execution of the interrupt processing.

[0079] When the interrupt processing is finished, a method "ExitFromInterrupt()" is executed to open the interrupt processing thread. As a result, a context switch is performed from the interrupt processing thread to the message processing thread, whereby the message processing B which has been suspended is restarted (resumed). When the message processing B is finished, a method "Exit()" is executed so as to open the message processing thread.

[0080] Thus, in the example shown in Fig. 4, a period exists in which the interrupt mask is open even when the message processing B is being executed. The message processing B is suspended when the interrupt mask is opened, so as to allow the interrupt to be executed. Thus, the illustrated embodiment enables the interrupt mask to be opened and closed in each period of message processing, by calling the method "BlockInterrupt()" and the message "AllowInterrupt()" during such a period. Namely, the interrupt handler is started to execute the interruption immediately after the method "AllowInterrupt()" is invoked, even during execution of message processing.

[0081] In this embodiment, therefore, an interrupt can be executed in response to an interrupt request without being held unexecuted until the end of the message processing, thus minimizing the interrupt latency. In addition, no specific consideration is necessary as to synchronization of the message processing with any processing other than the interrupt processing because the message processing is executed with a single thread. This greatly facilitates the programming work.

[0082] As has been fully described, in accordance with the type 2 programming, an interrupt processing thread is allocated as a thread which is used exclusively for the interrupt handler when the latter is registered. The arrangement is such that the period over which the interrupt mask is closed coincides with the period over which the interrupt processing thread runs. In other words, an interrupt is accepted whenever the interrupt processing thread is not working, even when the mes-

sage processing thread is working, so that the interrupt latency can be shortened.

[0083] It is also to be understood that the programming employed in this embodiment is not a mere multi-thread programming. Namely, the multi-thread programming employed in this embodiment is a specific one in which the message processing is executed by a single message processing thread alone. Therefore, the advantages which are brought about when the device drivers are described as parallel objects are maintained without being impaired.

[0084] The type 2 programming needs the exclusive control during the message processing only for the purpose of dealing with an interrupt conducted by the interrupt handler using the interrupt thread. Therefore, in this embodiment, the portion which undertakes the exclusive control closes itself in each device driver. In other words, the exclusive control which is conducted in each device driver for the purpose of achieving synchronization between the message processing and the interrupt processing does not affect the operation of other device drivers. This offers advantages such as implementation of a system that permits easy dynamic changes of device drivers. For instance, a system can be obtained in which a device driver can dynamically be changed by down-loading a program module through a network.

### 2-3-3 Interrupt mask in accordance with the embodiments

#### (Version 2)

[0085] Another embodiment of the present invention will now be described. This embodiment employs a programming style in which device drivers are described in the form of multi-thread objects and a message processing thread also is employed in an interrupt processing. This programming style will be referred to as "Type 3".

[0086] Interrupt processing is executed in an exclusive manner within a single interrupt thread. Therefore, the above-described programming style type 2 poses a problem in that, when consecutive interrupt messages "InterruptA" and "InterruptB" are received as shown in Fig. 5, execution of the interrupt processing corresponding to the later message "InterruptB" is held without being executed until execution of the interrupt processing responding to the earlier message "InterruptA" is finished. The programming style type 3 is intended to obviate this problem thereby further reducing the interrupt latency.

[0087] More specifically, in accordance with the type 3 programming style, an interrupt is divided into a plurality of sub-interrupts having different levels of priority. An interrupt processing thread is allocated only for the sub-interrupt of the top priority, and other sub-interrupts of lower levels of priority are assigned to a message processing thread. For instance, the top priority is given to a sub-interrupt that has to be executed without delay

for the purpose of controlling the hardware. Sub-interrupts of lower priority levels assigned to the message processing thread are, for example, those which are irrelevant to the control of the hardware, i.e., interrupts which need not be executed right away, such as copying of data into a common memory, reply to a client, and so forth.

[0088] A description will now be given of an execution model under the Type 3 programming style, with specific reference to Fig. 6 which assumes the same scenario as that described before.

[0089] Referring to Fig. 6, a message "RequestA" is input to the device driver, so that a message processing A is started in response to the message "RequestA" (invoked by message passing). The message processing A is executed by using a message processing thread. In the message processing A, a period shown by hatching, from the moment at which a method "BlockInterrupt()" is called to the moment at which a method "AllowInterrupt()" is called, is the period in which the message processing is executed with preference to the execution of an interrupt. Any interrupt received during this period, therefore, has to be kept waiting.

[0090] When the processing to be executed by the device driver is finished, the device driver executes a method "Exit()", regardless of the status or progress of the processing performed by the hardware, thereby opening the message processing thread. In this state, the device driver delays replying to the message "RequestA" until an interrupt indicating completion of the processing by the hardware is received from the hardware.

[0091] Then, a message "RequestB" is input to the device driver, so that a message processing B is started in response to the received message "RequestB" (invoked by message passing). The message processing B is executed by using the message processing thread. In this message processing B too, a period shown by hatching, from the moment at which a method "BlockInterrupt()" is called to the moment at which a method "AllowInterrupt()" is called, executes the message processing with preference to the execution of an interrupt. Any interrupt received during this period, therefore, has to be kept waiting.

[0092] It is assumed here that the processing which has been under execution by the hardware in response to the message "RequestA" is completed during the execution of the message processing B. This is an event that requests an interrupt to the device driver. Thus, a message "Interrupt", indicating a request for executing a predetermined processing which follows completion of the processing by the hardware, is input from the hardware to the device driver through the operating system.

[0093] The moment at which this interrupt has occurred falls within the period between the moment at which the method "BlockInterrupt()" is called and the moment at which the method "AllowInterrupt()" is called. Therefore, the interrupt message "Interrupt" is kept wait-

ing without being executed.

[0094] When the method "AllowInterrupt()" is called during the execution of the message processing B, the message processing B is suspended and a context switch is performed from the message processing thread to the interrupt processing thread, whereby the interrupt processing responding to the interrupt message "Interrupt" is started (invoked for interrupt handling). This interrupt processing is executed by means of the interrupt processing thread. Thus, in this embodiment, the interrupt latency corresponds to the period from the moment at which the interrupt has occurred to the moment at which the interrupt processing is enabled by the method "AllowInterrupt()".

[0095] The device driver allows the interrupt handler to execute the predetermined interrupt processing corresponding to the interrupt message "Interrupt". In the illustrated embodiment relying on the type 3 programming style, the interrupt handling enabled to execute the interrupt processing divides the interrupt processing into a portion, i.e., a sub-interrupt, which needs to be executed without delay and other portions which need not be executed right away. The sub-interrupt processing which has to be executed without delay is then executed by means of the interrupt processing thread, whereas other sub-interrupts which need not be executed right away are assigned to the message processing thread so as to be executed only after the completion of the message processing B, as will be explained later.

[0096] When the sub-interrupt processing which needs prompt execution is completed by the interrupt processing thread, a context switch is performed from the interrupt processing thread to the message processing thread, whereby the message processing B which has been suspended is restarted (resumed). When the message processing B is completed, the method "Exit()" is executed to open the message processing thread.

[0097] Thus, the message processing thread is now available for the processing of sub-interrupts which have been held without being executed as being unnecessary to execute right away. These sub-interrupts are therefore executed by means of the message processing thread. As a result of execution of these sub-interrupts, a reply (Reply for RequestA) is sent in response to the message "RequestA", whereby the whole interrupt processing is completed. Then, a method "ExitFromInterrupt()" is executed to open the message processing thread.

[0098] Thus, in accordance with the type 3 programming style, the interrupt processing is divided into a portion which needs a prompt execution and other portion or portions which do not require prompt execution. Only the portion of the processing which needs prompt execution is executed by means of the interrupt processing thread, whereas the other portion or portions of the interrupt processing are executed by means of the message processing thread. It is thus possible to further shorten the period over which the interrupt mask is kept

closed.

[0099] In the scheme shown in Fig. 6, the length of the interrupt latency is equivalent to that of the case shown in Fig. 4, although the period over which the interrupt mask is closed has been shortened. This is because only one interrupt has occurred during the execution of the message processing B.

[0100] A description will therefore be given as to an example in which the interrupt latency is appreciably shortened by virtue of the Type 3 of programming style, with specific reference to Fig. 7.

[0101] Referring to Fig. 7, a message "RequestA" is input to the device driver, and a message processing A is started in response to the message "RequestA" (invoked by message passing). The message processing A is executed by means of the message processing thread. When the processing to be conducted by the device driver is completed, the device driver finishes the message processing A to open the message processing thread, regardless of the status or progress of the processing conducted by the hardware. The device driver delays replying to the message "RequestA" until an interrupt indicating completion of the processing in response to the message "RequestA" is received from the hardware.

[0102] Then, another message "RequestB" is input to the device driver, so that a message processing B in response to the message "RequestB" is started (invoked by message passing). The message processing B is executed by means of the message processing thread. When the processing to be conducted by the device driver is completed, the device driver finishes the message processing B thereby opening the message thread, regardless of the status or progress of the processing conducted by the hardware. The device driver delays replying to the message "RequestB" until an interrupt indicating completion of the processing based on the message "RequestB" is received from the hardware.

[0103] Then, the processing which has been under execution by the hardware in response to the message "RequestA" is completed. As a result, an interrupt message "InterruptA" is input from the hardware to the device driver via the operating system requesting execution of a predetermined interrupt processing which follows the completion of the processing conducted by the hardware, whereby an interrupt processing A corresponding to the interrupt message "InterruptA" is started (invoked for handling interruptA). The interrupt handler divides the interrupt processing into a sub-interrupt which has to be executed without delay and a sub-interrupt which need not be executed right away. Only the sub-interrupt that needs prompt execution is executed without delay by means of the interrupt processing thread, whereas the sub-interrupt which does not require prompt execution is assigned to the message processing thread.

[0104] It is assumed that the processing which has

been under execution by the hardware responding to the message "RequestB" is completed so that an interrupt message "InterruptB" requesting an interrupt is input from the hardware to the device driver via the operating system while the sub-interrupt which needs prompt execution is being executed by the interrupt processing thread. At this moment, however, the interrupt is not accepted because the interrupt processing thread is working, so that the interrupt processing B in response to the interrupt message "InterruptB" from the hardware is held without being executed until the interrupt processing thread is opened.

[0105] The interrupt processing thread is opened when the processing of the portion, i.e., the sub-interrupt, of the interrupt processing A requiring prompt execution is finished. The interrupt processing B corresponding to the interrupt message "InterruptB" is then started by using the interrupt processing thread which is now open (invoked for handling interruptB).

[0106] It is to be understood that the interrupt processing B responding to the interrupt message "InterruptB" is started at the moment when the processing is finished for the portion, i.e., the sub-interrupt, of the interrupt processing A in response to the interrupt message "InterruptA" that requires execution without delay. Thus, the interrupt processing B is started without waiting for the completion of the whole interrupt processing A corresponding to the interrupt message "InterruptA". Therefore, the interrupt latency for the interrupt processing B corresponding to the interrupt message "InterruptB" is short as compared with the case of the type 2 programming style.

[0107] The interrupt processing B in response to the interrupt message "InterruptB" also is divided by the interrupt handler into a portion, i.e., a sub-interrupt that has to be executed without delay and an other portion, i.e., an other sub-interrupt, that does not need prompt execution. Only the sub-interrupt that needs prompt execution is executed without delay by means of the interrupt processing thread, whereas the sub-interrupt that need not be executed promptly is assigned to the message processing thread.

[0108] A context switch is performed from the interrupt processing thread to the message processing thread when the processing with the interrupt processing thread is completed, i.e., when the portion of the interrupt processing B corresponding to the interrupt message "InterruptB" that needs prompt execution is finished. Then, the portion of the interrupt processing A which has been kept unexecuted as being not urgent is executed by means of the message processing thread. When this processing is finished, the reply to the message "RequestA" which has been held pending is issued (Reply for RequestA). Then, the portion of the interrupt processing B which has been kept unexecuted as being not urgent is executed by means of the message processing thread. When this processing is finished, the reply to the message "RequestB" which has been held

pending is issued (Reply for RequestB).

[0109] As will be seen from the foregoing description, the type 3 programming style shortens the period over which the interrupt mask is closed, further to that under the type 2 programming style, thus achieving a further reduction in the interrupt latency.

#### 2-4 Operating environments of device drivers

[0110] Data processing method and apparatus have been described based on the types 2 and 3 programming styles, in comparison with the conventional technique which relies upon the type 1 programming style. It is preferred that the operation environments for these three types of programming style have upward compatibility, such that the device driver of the type 1 can operate in the environment of the type 2 and the environment of the type 3 permits device drivers of the types 1 and 2 to operate.

[0111] Such an upward compatibility is realized by providing an environment suited to a device driver that requires the greater number of functions than other device drivers. With such an upward compatibility, all the device drivers of different types of programming style in a system can operate without fail.

[0112] Device drivers of the type 2 requires a greater number of functions to be provided by the operating system than the device drivers of the type 1 because they are multi-threaded. Similarly, device drivers of the type 3 requires a greater number of functions to be provided by the operating system than the device drivers of the type 2, because under the type 3 an interrupt processing is divided into a portion to be executed in the interrupt processing thread and a portion to be executed in the message processing thread.

[0113] Consequently, the operating system required to provide environment suited to the type 2 has to have a greater size than the operating system which provides environment suited to the type 1. Similarly, the operating system required to provide environment suited to the type 3 has to have a greater size than the operating system which provides environment suited to the type 2.

[0114] This means that the use of an operating system capable of environment suited to the type 3 is not always preferred, although such an operating system can reduce the interrupt latency. Thus, the selection of the operating system subjects to factors including the degree of demand for reduction in the interrupt latency, as well as required functions, so as to avoid the use of excessively large operating system.

[0115] As will be seen from the foregoing description, according to the present invention, it is possible to reduce interrupt latency without impairing advantages brought about by the use of object-oriented device drivers.

[0116] In summary, embodiments of the present invention relate to a data processing method that implements device control in an object-oriented operating

system, and/or to a recording medium storing a program that implements that data processing method, as well as to a data processing apparatus incorporating the recording medium. Details of certain features referred to above are described in EP-A-0,753,811.

[0117] Although the invention has been described through its preferred forms, it is to be understood that the described embodiments are only illustrative and various changes and modifications may be imparted thereto without departing from the scope of the present invention which is limited solely by the appended claims.

#### Claims

##### 1. A data processing method, comprising:

describing a device driver for driving a hardware device in terms of a multi-thread object capable of allocating a message processing thread and at least one interrupt processing thread;  
allowing said device driver to assign to said message processing thread a processing based on a message received by said device driver from another object thereby executing said processing in said message processing thread; and  
allowing, when an event has occurred requesting an interrupt to said device driver, said device driver to execute the interrupt processing corresponding to said event in a processing thread corresponding to said event, wherein, in case that said corresponding interrupt processing thread is busy due to execution of another interrupt processing in response to another interrupt when said event has occurred, the interrupt processing corresponding to said event is held without being executed regardless of the state of said message processing thread and is executed only after completion of said another interrupt processing in said corresponding interrupt processing thread.

##### 2. A data processing method according to Claim 1, comprising:

allowing changeably setting an order of priorities between processing in said message processing thread and processing in said interrupt processing thread;  
wherein, in case that a processing given a higher priority than the interrupt processing is being executed in said message processing thread when an interrupt is received by said device driver, execution of the interrupt processing corresponding to the received interrupt is delayed until the processing in said message

processing thread is finished or until the setting of the priority order is changed so as to give higher priority to the interrupt processing than to the processing in said message processing thread, whereas, in case that the processing given the higher priority than the interrupt processing is not being executed in said message processing thread when said interrupt is received by said device driver, said interrupt processing is executed in said interrupt processing thread without delay after the receipt of the interrupt.

3. A data processing method according to Claim 1, wherein an interrupt processing is divided into a plurality of portions having different levels of priority, and a portion of the interrupt processing of a higher priority is executed by being assigned to said interrupt processing thread while a portion of a lower priority is executed by being assigned to said message processing thread.
4. A data processing method according to Claim 1, wherein a plurality of said device drivers are employed, and wherein said device drivers are described as parallel objects that can operate in parallel on an operating system.
5. A computer-readable recording medium storing a device driver which is a program for driving a hardware device,

wherein said device driver is described in terms of a multi-thread object capable of allocating a message processing thread and at least one interrupt processing thread;  
said device driver assigns to said message processing thread a processing based on a message received by said device driver from another object thereby executing said processing in said message processing thread; and  
wherein, when an event has occurred requesting an interrupt to said device driver, said device driver executes the interrupt processing corresponding to said event in a processing thread corresponding to said event, wherein, in case that said corresponding interrupt processing thread is busy due to execution of another interrupt processing in response to another interrupt when said event has occurred, the interrupt processing corresponding to said event is held without being executed regardless of the state of said message processing thread and is executed only after completion of said another interrupt processing in said corresponding interrupt processing thread.

6. A computer-readable recording medium according

to Claim 5,

wherein said device driver is arranged to allow setting of an order of priorities between processing in said message processing thread and processing in said interrupt processing thread, and

wherein, in case that a processing given a higher priority than the interrupt processing is being executed in said message processing thread when an interrupt is received by said device driver, execution of the interrupt processing corresponding to the received interrupt is delayed until the processing in said message processing thread is finished or until the setting of the priority order is changed so as to give higher priority to the interrupt processing than to the processing in said message processing thread, whereas, in case that the processing given the higher priority than the interrupt processing is not being executed in said message processing thread when said interrupt is received by said device driver, said interrupt processing is executed in said interrupt processing thread without delay after the receipt of the interrupt.

7. A computer-readable recording medium according to Claim 5, wherein said device driver divides an interrupt processing into a plurality of portions having different levels of priority, and a portion of the interrupt processing of a higher priority is executed by being assigned to said interrupt processing thread while a portion of a lower priority is executed by being assigned to said message processing thread.
8. A computer-readable recording medium according to Claim 5, wherein a plurality of said device drivers are employed, and wherein said device drivers are described as parallel objects that can operate in parallel on an operating system.
9. A data processing apparatus, comprising a recording medium storing a device driver which is a program for driving a hardware device and which is readable by said apparatus from said recording medium,

wherein said device driver is described in terms of a multi-thread object capable of allocating a message processing thread and at least one interrupt processing thread;  
said device driver assigns to said message processing thread a processing based on a message received by said device driver from another object thereby executing said processing in said message processing thread; and

wherein, when an event has occurred requesting an interrupt to said device driver, said device driver executes the interrupt processing corresponding to said event in a processing thread corresponding to said event, wherein, in case that said corresponding processing thread is busy due to execution of another interrupt processing in response to another interrupt when said event has occurred, the interrupt processing corresponding to said event is held without being executed regardless of the state of said message processing thread and is executed only after completion of said another interrupt processing in said corresponding interrupt processing thread.

10. A data processing apparatus according to Claim 9,

wherein said device driver is arranged to allow setting of an order of priorities between processing in said message processing thread and processing in said interrupt processing thread, and wherein, in case that a processing given a higher priority than the interrupt processing is being executed in said message processing thread when an interrupt is received by said device driver, execution of the interrupt processing corresponding to the received interrupt is delayed until the processing in said message processing thread is finished or until the setting of the priority order is changed so as to give higher priority to the interrupt processing than to the processing in said message processing thread, whereas, in case that the processing given the higher priority than the interrupt processing is not being executed in said message processing thread when said interrupt is received by said device driver, said interrupt processing is executed in said interrupt processing thread without delay after the receipt of the interrupt.

11. A data processing apparatus according to Claim 9, wherein said device driver divides an interrupt processing into a plurality of portions having different levels of priority, and a portion of the interrupt processing of a higher priority is executed by being assigned to said interrupt processing thread while a portion of a lower priority is executed by being assigned to said message processing thread.

12. A data processing apparatus according to Claim 9, wherein a plurality of said device drivers are employed, and wherein said device drivers are described as parallel objects that can operate in parallel on an operating system.

FIG. 1

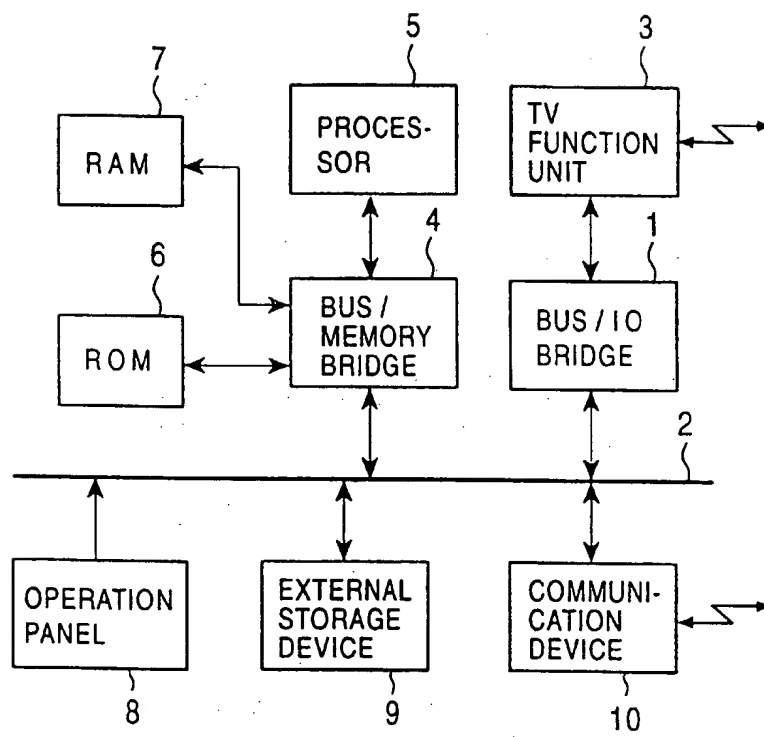


FIG. 2

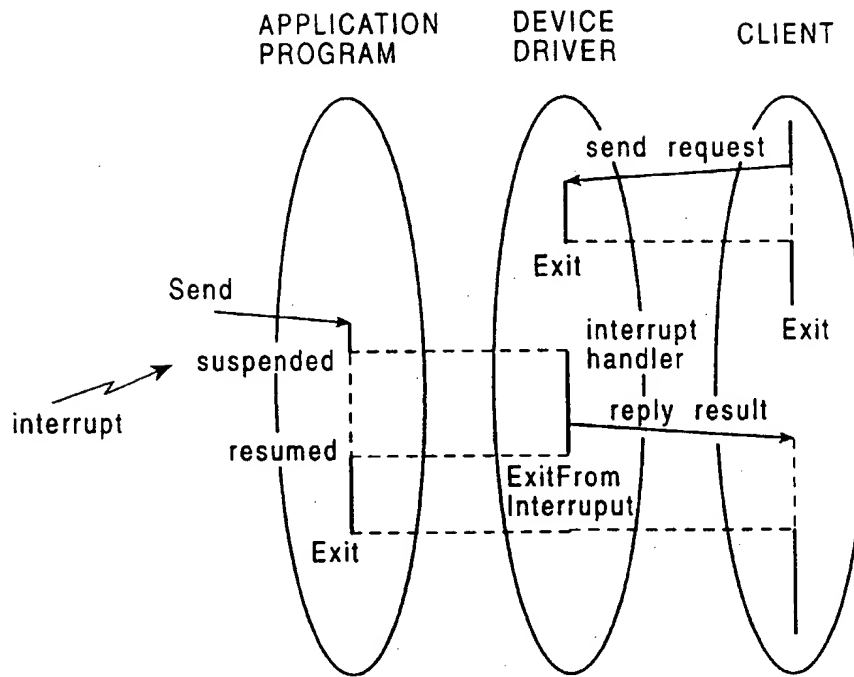




FIG. 3

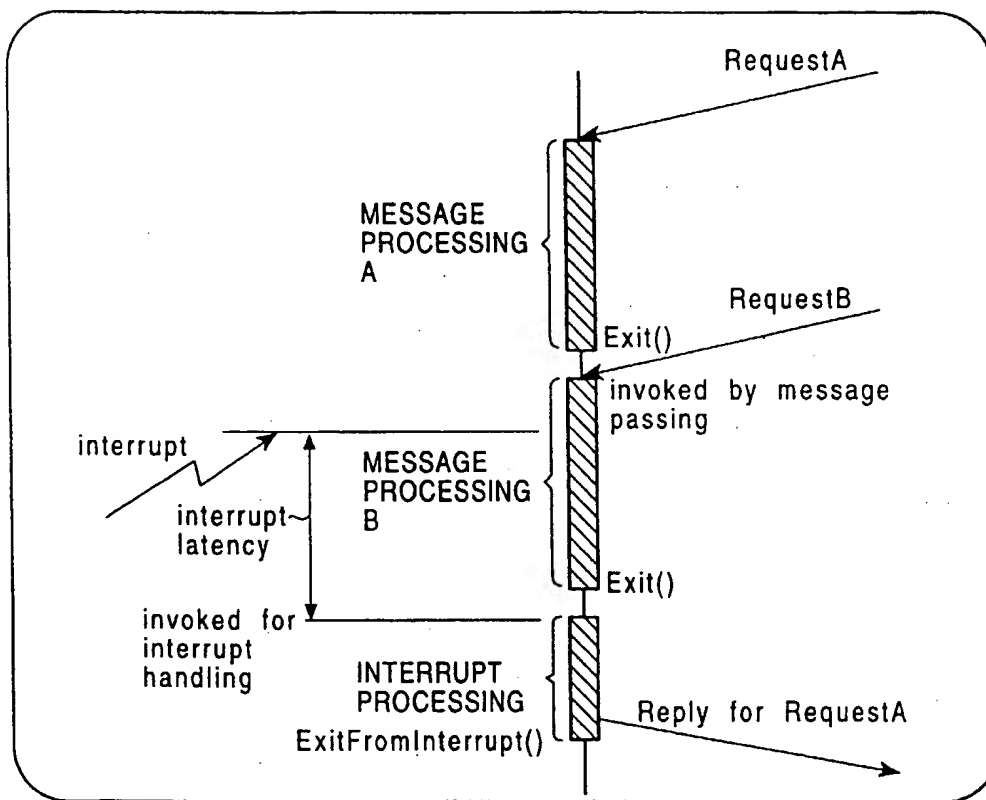


FIG. 4

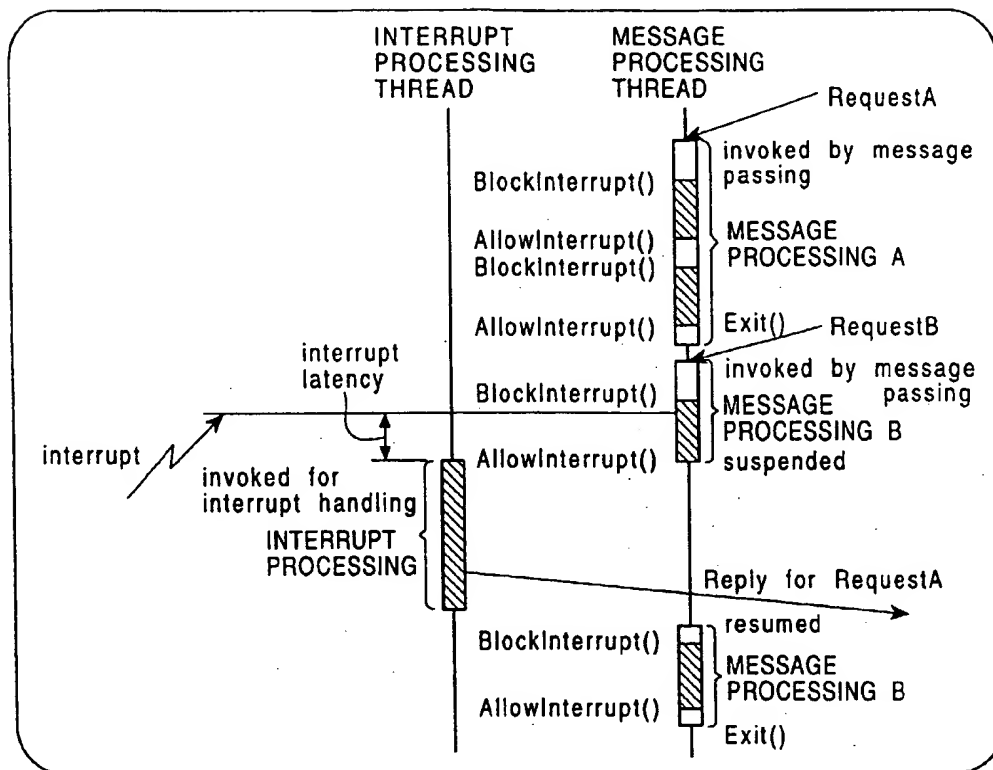


FIG. 5

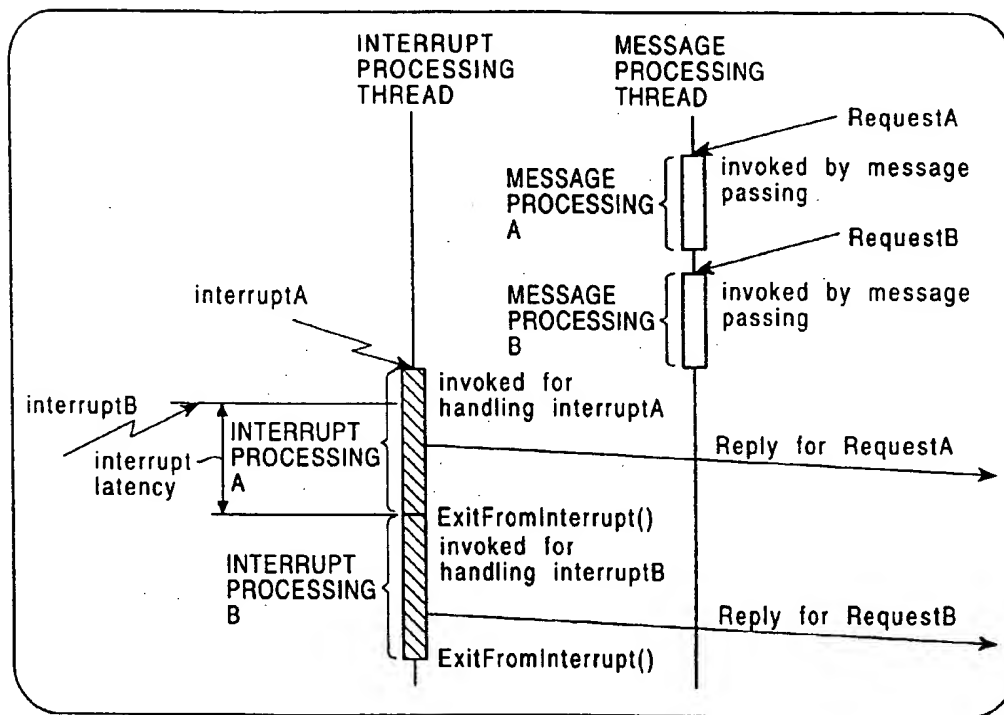


FIG. 6

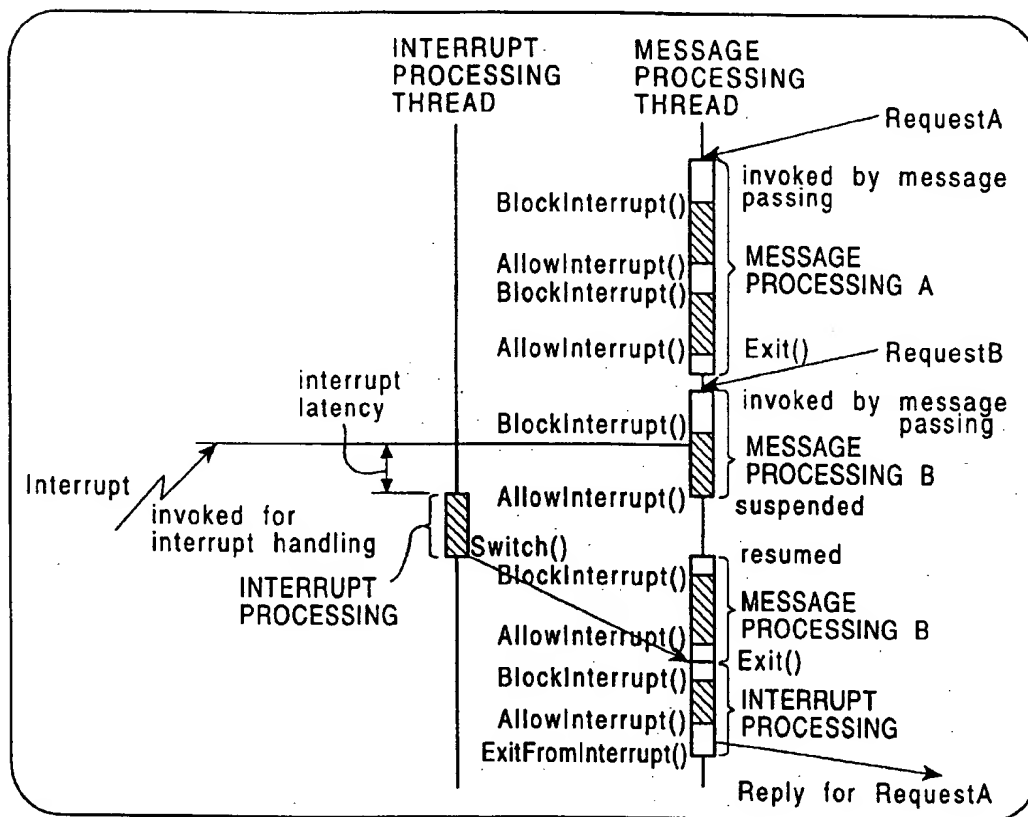


FIG. 7

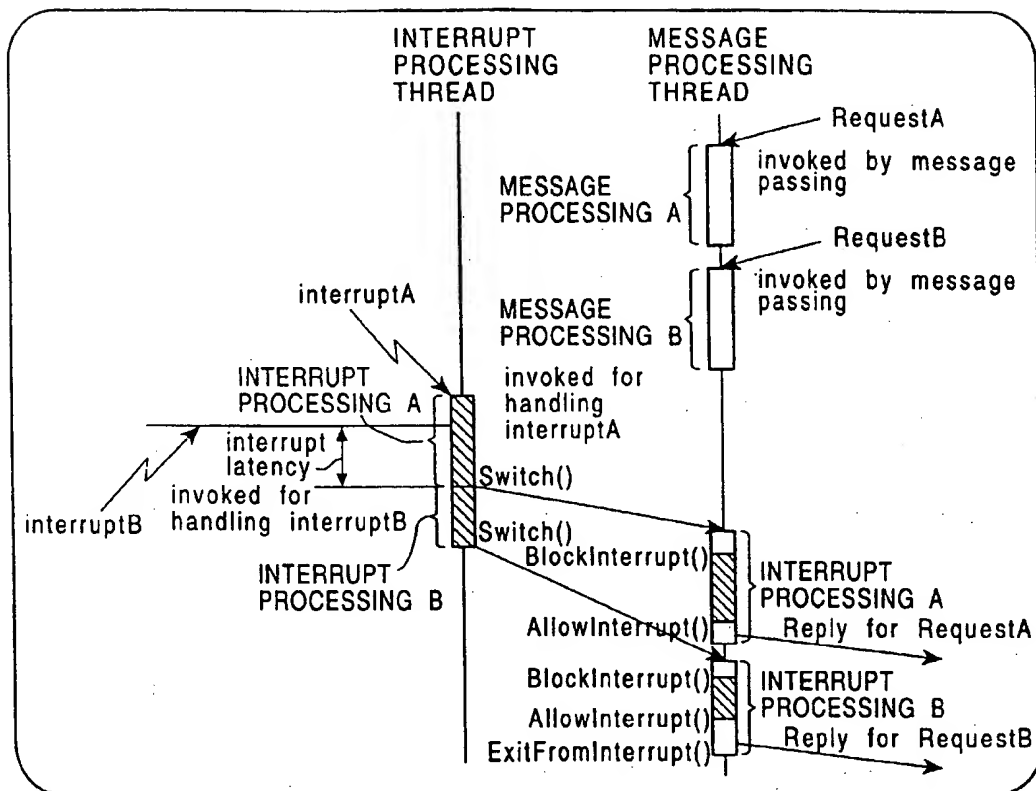


FIG. 8

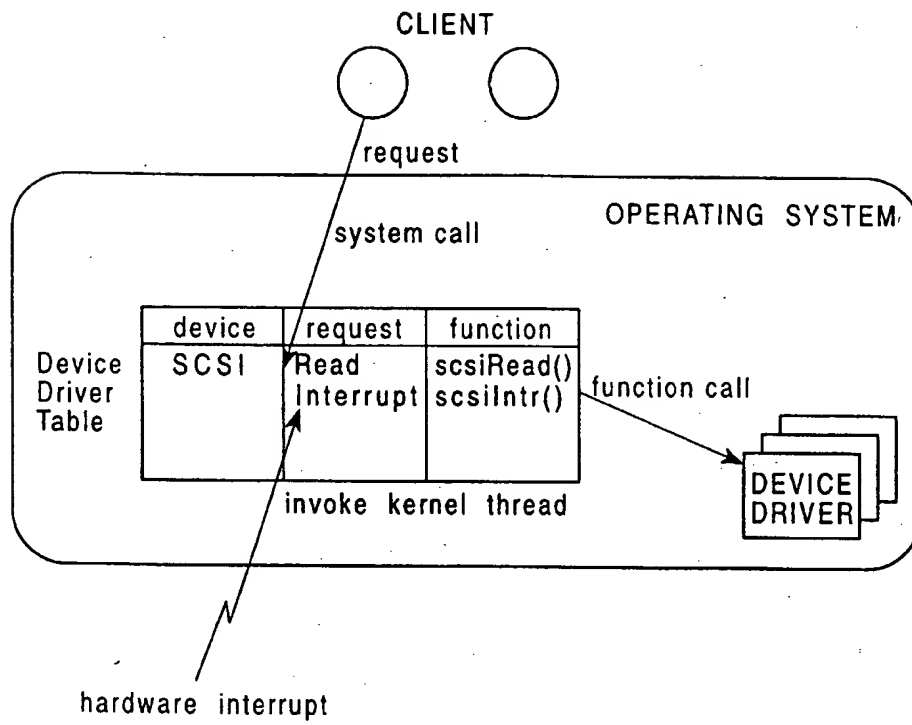
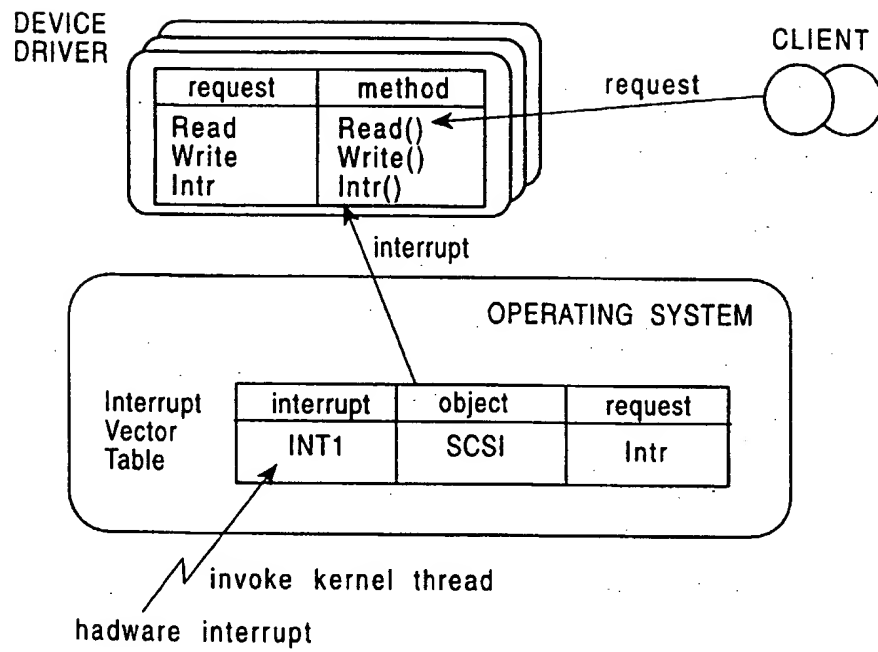


FIG. 9





European Patent  
Office

# EUROPEAN SEARCH REPORT

Application Number  
EP 98 30 6916

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.6)
A	WO 95 17719 A (TALIGENT INC) 29 June 1995 * the whole document *	1,4,9	G06F9/46 G06F13/10
A	LEMON S ET AL: "on object oriented device driver model" DIGEST OF PAPERS. COMPCON '95. TECHNOLOGIES FOR THE INFORMATION SUPERHIGHWAY (CAT. NO.95CH35737), DIGEST OF PAPERS. COMPCON'95. TECHNOLOGIES FOR THE INFORMATION SUPERHIGHWAY, SAN FRANCISCO, CA, USA, 5-9 MARCH 1995, pages 360-366, XP000566086 ISBN 0-8186-7029-0, 1995, Los Alamitos, CA, USA, IEEE Comput. Soc. Press, USA * page 363, left-hand column, line 33 - page 365, left-hand column, paragraph 3 *	1,4,9	
			TECHNICAL FIELDS SEARCHED (Int.Cl.6)
			G06F
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 2 December 1998	Examiner Fonderson, A
<p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document</p> <p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons &amp; : member of the same patent family, corresponding document</p>			

EPO FORM 1503 03/82 (P4/C01)